

Typical Solutions of Multi-User Linearly-Decomposable Distributed Computing

Ali Khalesi and Mohammad Reza Deylam Salehi

Abstract—We solve, in the typical-case sense, the multi-sender linearly-decomposable distributed computing problem introduced by tessellated distributed computing. We model real-valued encoders/decoders and demand matrices, and assess structural fidelity via a thresholded graph edit distance between the demand support and the two-hop support of the computed product. Our analysis yields: a closed-form second-moment (Frobenius) risk under spike-and-slab ensembles; deterministic links between thresholded GED and norm error; a Gaussian surrogate with sub-exponential tails that exposes explicit recall lines; concentration of GED and operator-norm control; and a compute-capped design with a visible knee. We map the rules to aeronautical and satellite networks.

Index Terms—Distributed computing, tessellation, GED, random products, aeronautical networking, satellite systems.

I. INTRODUCTION

Multi-user distributed computing (MUDC) increasingly runs at the edge and on airborne or satellite assets (UAVs, HAPS, LEO/GEO constellations), where both *communication* and *computation* budgets are tight, yet users demand low retrieval error. Classical approaches such as coded MapReduce, gradient coding, and coded matrix multiplication provide strong guarantees on norm error and straggler/latency resilience. However, many real failures are *structural*: a requested subfunction is simply unreachable through the server layer, even when the numerical error is small.

The tessellated framework [1] addresses the K -user, N -server distributed computing scenario, where each user requests a real function that is *linearly-decomposable* over L basis subfunctions. The set of user demands is represented by the coefficient matrix $F \in \mathbb{R}^{K \times L}$. Servers compute subsets of the L subfunctions (captured by the matrix E), while users connect to subsets of servers (captured by the matrix D). The key system parameters are: (i) the number of users K , which dictates demand dimensionality; (ii) the number of servers N , which captures available parallelism; and (iii) the number of subfunctions L , which reflects computational granularity. Performance is jointly determined by the reconstruction error ϵ (distortion), the computation load γ (fraction of subfunctions per server), and the communication load δ (fraction of users per server). Under zero-error operation, the problem reduces to finding sparse factors (D, E) such that $DE = F$ under strict per-row/column sparsity constraints. In the lossy case, the framework connects to approximate sparse matrix factorization, with achievable typical distortions characterized through

Marchenko–Pastur limits. This triplet (K, N, L) together with $(\gamma, \delta, \epsilon)$ provides the canonical tessellated system model, which our work extends by introducing thresholded graph-edit metrics for structural fidelity.

From norm error to structural fidelity: Motivated by deployments where reachability dominates, we complement the tessellated norm-centric picture with a structural metric: *thresholded graph edit distance (GED)* between the *demand support* of F and the *two-hop support* induced by (D, E) after a reliability test (thresholding of the real-valued two-hop strengths). This yields risk decompositions and regime maps in which GED bounds can dominate spectral/sub-multiplicative bounds in ultra-sparse two-hop settings, while ℓ_1/ℓ_2 risk dominates in moderate/dense regimes.

Positioning with fixed support factorization: The typical MUDC objective $\min_{D, E} \|DE - F\|_F^2$ under dimensionality and sparsity budgets is NP-hard in general and sits at the intersection of compressed sensing, low-rank approximation, and fixed-support matrix factorization. Unlike compressed sensing (which emphasizes uniqueness and recovery search), our design ties D to the instance F and budgets sparsity where it operationally matters (server fan-in/out), enabling capacity converses/achievability and explicit lossy laws.

a) Contributions: We provide a “typical solutions” framework for multi-user *linearly-decomposable* computing with real D, E, F that complements tessellated constructions:

- **Closed-form typical risk:** Exact $\mathbb{E}\|DE - F\|_F^2$ under independent spike-and-slab (real-valued) entries and a *thresholded* two-hop reach model q_τ that captures detection/reliability in real systems (Section II).
- **Thresholded-GED links:** Deterministic lower/upper connections between thresholded GED and Frobenius/ ℓ_1 risk, including when GED bounds beat spectral/Sub-multiplicative bounds (ultra-sparse two-hop) and when they do not (moderate/dense) (Sections III–IV).
- **Gaussian surrogate for reachability:** A tractable surrogate q_τ that exposes *recall lines* and *compute caps*, enabling one-pass operating-point selection (Section V).
- **Concentration and Boundary Rule:** Bounded-differences concentration for normalized GED and sub-exponential tail bounds for two-hop strengths under real D, E (Section VI), together with a simple slope test $s = c_+(1 - p_F) - c_-p_F$ that determines whether to maximize or minimize q_τ under a recall constraint, thereby unifying budget and SLA tuning.
- **Practical guidance:** *Aeronautical* and *satellite* deployments: mapping δ, γ to link and compute caps; accommodating link intermittency, Doppler, and handover; and

A. Khalesi is an *Assistant Professor* at Institut Polytechnique des Sciences Avancées (IPSA), Paris, France (ali.khalesi@ipsa.fr). M.R. Deylam Salehi is with the Communication Systems Department, EURECOM, Biot Sophia Antipolis, France (deylam@eurecom.fr).

sizing thresholds to meet per-user recall while minimizing extra-fetch.

b) *Why multi-user linearly-decomposable is the right approach:* It strictly generalizes gradient coding and coded linear transforms by allowing heterogeneous user weights and multi-shot designs while supporting both lossless corner points and lossy typical laws—precisely the dual view (capacity vs. distortion) that tessellation champions. Our structural GED analysis fits naturally on top of these real-valued factorizations and explains when “coverage” (reachability) is the true bottleneck and how to buy it with minimal δ, γ .

II. SETUP AND THRESHOLDED GED

We consider K users, N servers, and L sub-functions. The demand matrix is $F \in \mathbb{R}^{K \times L}$, the user \rightarrow server encoding is $D \in \mathbb{R}^{K \times N}$, the server \rightarrow sub-function decoding is $E \in \mathbb{R}^{N \times L}$, and the arithmetic two-hop map is $A = DE \in \mathbb{R}^{K \times L}$. Entry $A_{k\ell}$ aggregates all two-hop contributions from user k to sub-function ℓ through the N servers.

Budgets and degrees: The sparsity budgets are the *communication density* p_D (fraction of nonzeros per row of D) and the *compute density* p_E (fraction of nonzeros per column of E). Expected degrees are $\mathbb{E}[\deg_{\text{user}}] = Np_D$ and $\mathbb{E}[\deg_{\text{subfn}}] = Np_E$. In networks, p_D maps to scheduling/fan-out on user \rightarrow server links; p_E maps to compute/I/O activation on server \rightarrow sub-function links.

Thresholded supports: Real systems read out connections through a detector or a reliability test. Let $\tau > 0$ be a decision threshold on two-hop strengths and let $\tau_F > 0$ (often $\tau_F = \tau/2$) be a demand relevance threshold. Define the induced supports

$$B = \mathbb{1}_{[|A| \geq \tau]} \in \{0, 1\}^{K \times L}, \quad S = \mathbb{1}_{[|F| \geq \tau_F]} \in \{0, 1\}^{K \times L}.$$

With unit edge costs, the thresholded graph edit distance (GED) between S and B is

$$\begin{aligned} \text{GED}_\tau(S, B) &= \sum_{k, \ell} \mathbb{1}_{[S_{k\ell}=1, B_{k\ell}=0]} \\ &\quad + \sum_{k, \ell} \mathbb{1}_{[S_{k\ell}=0, B_{k\ell}=1]}. \end{aligned} \quad (1)$$

It counts *misses* (false negatives) and *extras* (false positives) after thresholding. We also track the Frobenius error $\|A - F\|_F$ and the weighted edit (entrywise ℓ_1) $\|A - F\|_1$.

Spike-and-slab ensemble: To capture sparsity with real weights, we adopt independent spike-and-slab entries:

$$D_{kn} \sim (1 - p_D) \delta_0 + p_D \mathcal{N}(0, v_D), \quad (2)$$

$$E_{n\ell} \sim (1 - p_E) \delta_0 + p_E \mathcal{N}(0, v_E), \quad (3)$$

and similarly $F_{k\ell} \sim (1 - p_F) \delta_0 + p_F \mathcal{N}(0, v_F)$. Entries are independent across matrices. Under zero-mean slabs,

$$A_{k\ell} = \sum_{n=1}^N D_{kn} E_{n\ell}, \quad (4)$$

$$\mathbb{E}[A_{k\ell}] = 0, \quad \text{Var}(A_{k\ell}) = N p_D p_E v_D v_E. \quad (5)$$

Thus $A_{k\ell}$ is sub-exponential; by CLT it is well-approximated as Gaussian near the operating point, with standard deviation $\sigma_A = \sqrt{N p_D p_E v_D v_E}$.

Reachability, recall, and precision: Let $q_\tau = \mathbb{P}(|A_{k\ell}| \geq \tau)$ be the *two-hop reach probability* at threshold τ . Using the Gaussian surrogate,

$$q_\tau(p_D, p_E) \approx 2Q\left(\frac{\tau}{\sigma_A}\right), \quad (6)$$

where Q is the standard normal tail. The demand-support probability is

$$p_S = \mathbb{P}(|F_{k\ell}| \geq \tau_F) = (1 - p_F) \cdot 0 + p_F \cdot 2Q\left(\frac{\tau_F}{\sqrt{v_F}}\right) \quad (7)$$

With asymmetric costs (c_- , c_+) for misses/extras, the expected per-edge cost is the affine function

$$\phi_\tau(q_\tau) = c_- p_S (1 - q_\tau) + c_+ (1 - p_S) q_\tau. \quad (8)$$

Its slope $s = c_+ (1 - p_S) - c_- p_S$ determines whether to raise or lower q_τ (boundary rule).

Threshold selection and invariances: Scaling $(D, E) \mapsto (\alpha D, \alpha^{-1} E)$ leaves A unchanged; hence τ should be interpreted relative to scale, e.g., via the *dimensionless* threshold $\tilde{\tau} = \tau/\sigma_A$. Larger τ reduces extras (better precision) but increases misses (worse recall). In practice:

- choose τ from an ROC target (application SNR or minimum precision);
- set τ_F as the minimum “useful” magnitude in F (e.g., coefficient quantization or SLA threshold);
- enforce a recall requirement $\mathbb{P}(B=1 | S=1) \geq \rho$, which is met by $q_\tau \geq \rho$ under independence.

a) *Why thresholded GED:* In ultra-sparse two-hop regimes (small $p_D p_E$ or small N), entries of A concentrate near 0 or a single-path contribution. Then the structural decision “is the edge *there*?” dominates, and GED_τ tightly tracks the operational loss. As multiplicities grow (moderate/dense two-hop), magnitude errors matter and $\|A - F\|_1 / \|A - F\|_F$ become more faithful; our later bounds translate performance between these metrics.

b) *Network mapping:* On real links, p_D reflects per-slot scheduling or beam fan-out; p_E reflects server occupancy or accelerator duty cycle; τ encodes the demod/decoding or estimation reliability required for an edge to be “usable.” The triplet (p_D, p_E, τ) therefore exposes a concrete *communication-computation-retrieval* design plane, and GED_τ measures the structural gap induced by design choices.

III. DETERMINISTIC INEQUALITIES

We compile Frobenius- and operator-norm inequalities that hold for *any* real D, E, F , and discuss the conditions under which they are informative for design.

Note 1 (Reverse/forward triangle).

$$\left| \|DE\|_F - \|F\|_F \right| \leq \|DE - F\|_F \leq \|DE\|_F + \|F\|_F. \quad (9)$$

Proof. Apply the reverse and forward triangle inequalities with $X=DE$, $Y=F$ in the Hilbert space $(\mathbb{R}^{K \times L}, \langle \cdot, \cdot \rangle_F)$, where $\langle X, Y \rangle_F = \text{Tr}(X^\top Y)$. \square

Equality/tightness: The *upper* bound becomes equality when DE and F are positively colinear in Frobenius inner product, i.e., $\langle DE, F \rangle_F = \|DE\|_F \|F\|_F$. The *lower* (reverse)

bound becomes equality when they are negatively colinear. If $\langle DE, F \rangle_F = 0$ (Frobenius orthogonality), then

$$\|DE - F\|_F^2 = \|DE\|_F^2 + \|F\|_F^2, \quad (10)$$

which is a Pythagorean case useful for sanity checks.

Remark 1 (Sub-multiplicative upper bounds).

$$\|DE - F\|_F \leq \|DE\|_F + \|F\|_F \leq \|D\|_F \|E\|_2 + \|F\|_F,$$

$$\|DE - F\|_F \leq \|DE\|_F + \|F\|_F \leq \|D\|_2 \|E\|_F + \|F\|_F.$$

Proof. The first inequality is Note 1. For the second step use the standard submultiplicativity $\|AB\|_F \leq \|A\|_F \|B\|_2$ and symmetrically $\|AB\|_F \leq \|A\|_2 \|B\|_F$. \square

When is Remark 1 informative? The bound $\|DE\|_F \leq \|D\|_F \|E\|_2$ is *tight* when E acts nearly as an isometry on the dominant right singular subspace of D (e.g., E has a large top singular value aligned with D 's top right singular vectors). The symmetric form is tight when D has a large top singular value aligned with E 's top left singular vectors. In practice:

- If E contains high-degree “hubs” (large operator norm), $\|E\|_2$ drives the upper bound and warns of possible overestimation by purely spectral arguments (GED bounds may be better in ultra-sparse two-hop).
- If both $\|D\|_2$ and $\|E\|_2$ are modest but $\|D\|_F, \|E\|_F$ are large, the weaker $\|DE\|_F \leq \|D\|_F \|E\|_F$ (implied by $\|E\|_2 \leq \|E\|_F$) can still be numerically useful as a coarse ceiling.

Proposition 1 (Singular-value lower bounds).

$$\|DE - F\|_F \geq \left| \|D\|_F \sigma_{\min}(E) - \|F\|_F \right|, \quad (11)$$

$$\|DE - F\|_F \geq \left| \sigma_{\min}(D) \|E\|_F - \|F\|_F \right|. \quad (12)$$

Proof. We have $\|DE\|_F \geq \sigma_{\min}(E) \|D\|_F$ because

$$\begin{aligned} \|DE\|_F^2 &= \text{Tr}(E^\top D^\top DE) \geq \sigma_{\min}(E)^2 \text{Tr}(D^\top D) \\ &= \sigma_{\min}(E)^2 \|D\|_F^2, \end{aligned} \quad (13)$$

using the Courant–Fischer characterization on $E^\top D^\top DE$. Combine with the reverse triangle bound in Note 1. The symmetric inequality follows by swapping D and E . \square

Meaning for design: Prop. 1 says DE cannot be too small in Frobenius norm if either factor is well-conditioned (large σ_{\min}) while the other carries energy. Thus even if F is small, strong conditioning in E (or D) forces a nontrivial lower bound on the error unless the other factor is correspondingly small. This is a warning against *over-activating* well-conditioned blocks when the target F is structurally sparse or low power.

A quick orthogonality corollary.: If DE and F are orthogonal, then by the proof of Note 1

$$\begin{aligned} \|DE - F\|_F &= \sqrt{\|DE\|_F^2 + \|F\|_F^2} \geq \\ &\left| \|D\|_F \sigma_{\min}(E) - \|F\|_F \right|, \end{aligned} \quad (14)$$

which shows the singular-value lower bound is often close to exact in “energy-separated” regimes (useful when F focuses on entries not energized by DE).

Link to thresholded GED: Upper bounds like Remark 1 can be pessimistic in ultra-sparse two-hop regimes because $\|E\|_2$ is driven by hubs that do not change *support*. In such cases, the structural bounds GED_τ (and its links to $\|A - F\|_F$ through threshold factors) give sharper guidance for communication/compute budgeting; see Sections on *Norm Risk and Thresholded GED* and *Reach Probability, Recall Lines, and Cap*.

IV. NORM RISK AND THRESHOLDED GED

This section connects the squared Frobenius risk to thresholded GED via simple, interpretable decompositions. Throughout, $A = DE$, and entries follow the independent spike-and-slab model with zero-mean Gaussian slabs unless stated otherwise.

Lemma 1 (Second moment of $\|DE - F\|_F$). *Under independence and spike-and-slab with variances (v_D, v_E, v_F) and densities (p_D, p_E, p_F) , for any slab mean $\mu_F = \mathbb{E}[F_{k\ell}]$,*

$$\begin{aligned} \mathbb{E}\|DE - F\|_F^2 &= KL \left(N p_D p_E v_D v_E + p_F v_F \right. \\ &\quad \left. + (N p_D p_E v_D v_E - p_F \mu_F)^2 \right). \end{aligned} \quad (15)$$

With zero-mean slabs for F ($\mu_F = 0$), this reduces to

$$\begin{aligned} \mathbb{E}\|DE - F\|_F^2 &= KL \left(N p_D p_E v_D v_E + p_F v_F \right. \\ &\quad \left. + (N p_D p_E v_D v_E)^2 \right). \end{aligned} \quad (16)$$

Proof. For fixed (k, ℓ) , $A_{k\ell} = \sum_{n=1}^N D_{kn} E_{n\ell}$. Independence and zero-mean slabs for D, E give $\mathbb{E}[A_{k\ell}] = 0$ and $\text{Var}(A_{k\ell}) = \sum_n \mathbb{E}[D_{kn}^2] \mathbb{E}[E_{n\ell}^2] = N p_D v_D p_E v_E$. Then, using $\mathbb{E}[(A_{k\ell} - F_{k\ell})^2] = \text{Var}(A_{k\ell}) + \text{Var}(F_{k\ell}) + (\mathbb{E}[A_{k\ell}] - \mathbb{E}[F_{k\ell}])^2$, sum over KL entries to obtain (15). If $\mu_F = 0$ the third term becomes $(N p_D p_E v_D v_E)^2$. \square

Interpretation: The three terms in (15) have clear roles: (i) $N p_D p_E v_D v_E$ is the two-hop *variance* induced by DE (communication \times compute budget); (ii) $p_F v_F$ is the intrinsic *variance* of the demand; (iii) $(N p_D p_E v_D v_E - p_F \mu_F)^2$ is a *bias* term capturing mean-energy mismatch between A and F . In the zero-mean case (16), the “bias” reduces to the square of the two-hop energy.

Lemma 2 (Thresholded GED vs. Norm). *Let $B = \mathbb{1}_{\|DE\| \geq \tau}$ and $S = \mathbb{1}_{\|F\| \geq \tau_F}$. Then entrywise*

$$\left(\|DE\| - \|F\| \right)^2 \geq \min\{\tau, \tau_F\}^2 \mathbb{1}_{[B \neq S]}, \quad (17)$$

so

$$\text{GED}_\tau(S, B) \leq \frac{\| \|DE\| - \|F\| \|_F^2}{\min\{\tau, \tau_F\}^2} \leq \frac{\|DE - F\|_F^2}{\min\{\tau, \tau_F\}^2}. \quad (18)$$

Proof. If $B \neq S$, one of $\|DE\|$ or $\|F\|$ is \geq its threshold while the other is $<$ its threshold; hence $\| \|DE\| - \|F\| \| \geq \min\{\tau, \tau_F\}$. Summing squares across entries yields the first inequality; the second follows from $\|X\| - \|Y\| \leq \|X - Y\|$ entrywise. \square

Normalization remark: Many designs absorb the $\min\{\tau, \tau_F\}^2$ factor into unit costs (c_-, c_+) or fix $\tau_F = \tau$, recovering the simplified form $\text{GED}_\tau(S, B) \leq \| \|DE\| - \|F\| \|_F^2 \leq \|DE - F\|_F^2$ used for quick comparisons.

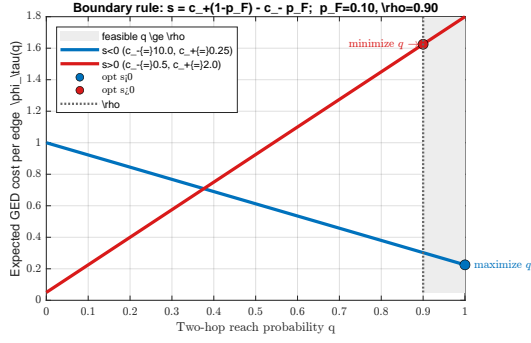


Fig. 1. Boundary rule: if $s < 0$ maximize q ; if $s > 0$ minimize q subject to $q \geq \rho$.

Lemma 3 (Weighted edit). *For any matrix X ,*

$$\frac{\|X\|_1}{\sqrt{KL}} \leq \|X\|_F \leq \|X\|_1, \quad (19)$$

$$\frac{\|DE-F\|_1}{\sqrt{KL}} \leq \|DE-F\|_F \leq \|DE-F\|_1. \quad (20)$$

Proof. Standard ℓ_1 - ℓ_2 norm relations in finite dimensions applied entrywise. \square

Why these links matter?

- *Sparse two-hop:* Most $A_{k\ell}$ are near zero or come from a single path; thresholding is the decision of record. Then GED_τ is a tight proxy, and (18) lower-bounds the norm risk up to a fixed (threshold) factor.
- *Moderate/dense two-hop:* Multiplicities $S_{k\ell} \geq 2$ are common; magnitudes matter. The ℓ_1 risk is more faithful than Hamming, while $\|DE-F\|_F$ captures energy mismatch as in Lemma 1.
- *Design dial:* The slope of the expected GED cost $\phi_\tau(q) = c_-p_S(1-q) + c_+(1-p_S)q$ decides whether to push q up or down. Because q is monotone in (p_D, p_E) , one obtains a one-pass policy along recall lines until a compute cap or network constraint binds (see (6)).

V. DESIGN MAP, EVALUATION, AND DOMAIN GUIDANCE

Boundary rule: The expected per-edge GED $\phi_\tau(q)$ is affine in q with slope $s = c_+(1 - p_F) - c_-p_F$. If $s < 0$ (misses costlier), *maximize* q (push p_D to caps, raise p_E up to the recall line/cap). If $s > 0$ (extras costlier), *minimize* q subject to the recall constraint.

Setup: We consider a distributed computation setting with $K = 64$ users, $N = 64$ servers, and $L = 800$ sub-functions, where the system parameters are $p_F = 0.10$, $(c_-, c_+) = (10, 0.25)$, $v_D = v_E = 0.50$, $\tau = 0.10$, $\tau_F = \tau/2$, $p_E^{\text{cap}} = 0.20$, and $\rho \in \{0.80, 0.90, 0.95\}$.

Aeronautical networks: In A2G/A2A, p_D maps to sector and beam scheduling and p_E to onboard/edge compute activation. If $s < 0$, increase p_D (coverage first), then lift p_E along the recall line; if $s > 0$, restrain p_D until recall binds to avoid extra fetch/traffic.

Satellite systems: In multi-beam LEO/GEO, p_D controls feeder fan-out and p_E the payload occupancy. The knee gives

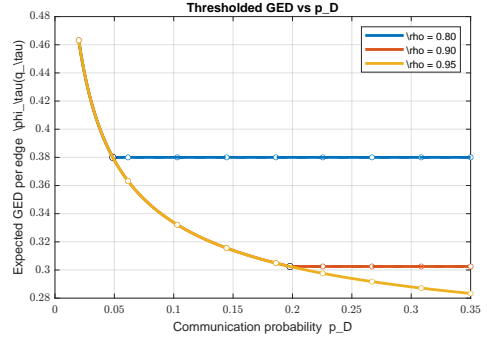


Fig. 2. Thresholded GED vs. p_D (solid: cap-binding; dashed: recall-binding). Circles mark the knee where $p_E^{\min}(p_D, \rho) = p_E^{\text{cap}}$.

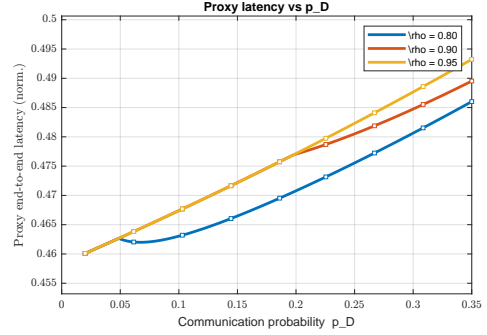


Fig. 3. Proxy latency vs. p_D under the same settings.

the smallest p_D that achieves the target recall at the compute cap, guiding beam hopping, task placement, and SLA tuning.

Fig. 1 illustrates the boundary rule, showing how the slope s dictates whether to maximize or minimize q . Fig. 2 plots the thresholded GED versus p_D , with the knee marking the smallest p_D that achieves the recall target at the compute cap. Fig. 3 reports the corresponding proxy latency trends under the same settings.

VI. CONCLUSION

We provided a typical-case solution for real-valued MUDC with thresholded GED: an affine objective in the reach probability, explicit recall lines, compute caps with visible knees, and concentration. Upper bounds from sub-multiplicativity dominate in dense regimes; GED-based bounds are tight and informative in sparse two-hop regimes. The slope-driven boundary rule offers a deployable design knob for aeronautical and satellite networks.

REFERENCES

- [1] A. Khalesi and P. Elia, “Tessellated Distributed Computing,” *IEEE Trans. Inf. Theory*, vol. 71, no. 6, pp. 4754–4784, Jun. 2025. doi:10.1109/TIT.2025.3556384
- [2] K. Wan, H. Sun, M. Ji, and G. Caire, “Distributed linearly separable computation,” *IEEE Trans. Inf. Theory*, vol. 68, no. 2, pp. 1259–1278, 2022.
- [3] S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, “Coded MapReduce,” in *Proc. Allerton*, 2015.
- [4] R. Tandon, Q. Lei, A. G. Dimakis, and N. Karampatziakis, “Gradient coding: Avoiding stragglers in distributed learning,” in *Proc. ICML*, 2017.